



US006477595B1

(12) **United States Patent**
Cohen et al.

(10) **Patent No.:** US 6,477,595 B1
(45) **Date of Patent:** Nov. 5, 2002

(54) **SCALABLE DSL ACCESS MULTIPLEXER WITH HIGH RELIABILITY**

(75) Inventors: Michael S. Cohen, Enola, PA (US); K. C. Babu, Bangalore (IN); R. Balasubramanyam, Bangalore (IN); Murali Ramanathan, Bangalore (IN); P. K. Nanda Kumar, Bangalore (IN)

(73) Assignee: E-Cell Technologies, Austin, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: 09/547,921

(22) Filed: Apr. 11, 2000

Related U.S. Application Data

(60) Provisional application No. 60/161,420, filed on Oct. 25, 1999.

(51) Int. Cl. 7 G06F 13/42

(52) U.S. Cl. 710/105; 710/104; 379/93.14;
379/93.28

(58) Field of Search 379/93.14, 93.28;
710/104, 105

(56) **References Cited****U.S. PATENT DOCUMENTS**

- RE31,182 E * 3/1983 Crager et al.
4,494,232 A * 1/1985 Dambrackas et al.
4,782,512 A * 11/1988 Hutton
5,062,104 A * 10/1991 Lubarsky et al.
5,202,884 A * 4/1993 Gaddis et al.
5,272,697 A * 12/1993 Fraser et al.
5,327,484 A * 7/1994 Connell et al.
5,345,437 A * 9/1994 Ogawa
5,471,472 A * 11/1995 McClure et al.
5,479,447 A * 12/1995 Chow et al. 375/260
5,513,251 A * 4/1996 Rockkind et al. 379/93.35

- 5,528,585 A * 6/1996 Cooley et al. 370/347
5,541,919 A * 7/1996 Yong et al.
5,546,379 A * 8/1996 Thaweechai et al. 370/254
5,649,217 A * 7/1997 Yamanaka et al.
5,657,314 A * 8/1997 McClure et al.
5,668,814 A * 9/1997 Balatoni
5,668,857 A * 9/1997 McHale
5,756,280 A * 5/1998 Soora et al.
5,781,617 A * 7/1998 McHale et al.
5,790,548 A * 8/1998 Sistanizadeh et al.
5,812,786 A * 9/1998 Seazholtz et al.
5,852,655 A * 12/1998 McHale et al. 379/93.14
5,898,761 A * 4/1999 McHale et al.
5,905,781 A * 8/1999 McHale et al.
5,936,963 A * 8/1999 Saussy
5,987,009 A * 11/1999 Awazu
5,999,565 A * 12/1999 Locklear, Jr., et al.
6,012,113 A * 1/2000 Tuckner
6,014,431 A * 1/2000 McHale et al.
6,029,197 A * 2/2000 Erimli et al.
6,061,357 A * 5/2000 Timm et al.
6,088,430 A * 7/2000 McHale

OTHER PUBLICATIONS

T. Jessup, "DSL: The Corporate Connection", Feb. 1, 1998, pp. 103-110.

PCT Int'l Search Report (7 pp).

* cited by examiner


Primary Examiner—Rupal Dharia


(74) *Attorney, Agent, or Firm*—McDermott, Will & Emery

(57) **ABSTRACT**

A digital subscriber line (DSL) multiplexer is used in conjunction with a DSL modem to provide a reliable high speed connection for an end-user. The scalable DSL multiplexer supports multiple DSL modems to a number of subscriber desiring increased bandwidth. The connection is reliably maintained by the multiplexer which utilizes a proprietary protocol to exchange management information to monitor the connection.

9 Claims, 9 Drawing Sheets








FIG. 1B
(PRIOR ART)




FIG. 2A




FIG. 2B




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8

SCALABLE DSL ACCESS MULTIPLEXER WITH HIGH RELIABILITY

This application claims the benefit of U.S. Provisional Application No. 60/161,420, filed Oct. 25, 1999, which is fully incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Technical Field

The present invention relates to a digital subscriber line (DSL) multiplexer for providing connectivity to a data network. In particular, the present invention involves the exchange of management information from the DSL multiplexer to an end-user DSL modem.

2. Background Art

As the information age matures, it is enabled by a number of technological advances, such as the geometric growth of networked computing power and the prevalence of reliable and ubiquitous transmission media. Today's consumers in both the residential and business arena have been acclimated to a more graphical approach to communication. In particular, multimedia applications (which include textual, graphical, image, video, voice and audio information) have become increasingly popular and find usage in science, business, and entertainment. Local area networks (LANs) are essential to the productivity of the modem workplace; Ethernet-type networks have dominated the LAN market and have been continually enhanced (e.g., switched Ethernet, Fast Ethernet, and/or Gigabit Ethernet) to keep pace with the bandwidth intensive multimedia applications.

A compelling example of the growth of information consumption is the dramatic increase in users of the World Wide Web, a multimedia-based information service provided via the Internet. Although initially a forum for academia to exchange ideas captured in ASCII text, the Internet has developed to become a global media for users from all walks of life. These Internet users regularly exchange multimedia graphical, image, video, voice and audio information as well as text.

Furthermore, the business world has come to realize tremendous value in encouraging workers to telecommute. To avoid the idle commuting time, today's workers enjoy the convenience of working from home via their personal computers. As illustrated in FIG. 1A, a user at a remote site 101 (e.g., home) has traditionally been able to access her/his office 119, which includes accessing an office local area network 119b (LAN), through a dial-up connection over a 33 Kbps or 56 Kbps modem 101b. The dial-up connection is handled by a telephone communication facility 240 (CO) 105 through a voice switch 107, which switches the "data" call through a public switched telephone network (PSTN) 111. The data call terminates in a remote CO 121 at a voice switch 123. The voice switch 123 switches the call to the subscriber; in this case, the called line is associated with a modem in a modem pool 119a. Once connected to the modem pool 119a, the end user at her/his remote site 101 can access the computing resources in his office 119. These sources include a multimedia server 119c and a PC 119d of the remote user. A similar connection to Internet 115 by a user at a remote site 101 can be accomplished by connecting to an Internet Service Provider (ISP) 117 instead of modem pool 119c.

Unfortunately, telecommuting from a remote office or accessing multimedia information from home over the Internet imposes an enormous strain on networking resources. It is common knowledge that the networking infrastructure is

the bottleneck to the expedient transfer of information, especially bandwidth intensive multimedia data. As alluded to before, today's access methods are limited to standard analog modems, such as 101b, which have a maximum throughput of 56 Kbps on a clean line (i.e., a line not having any appreciable noise causing errors in bit rate transfer). Remote users may alternatively acquire basic rate (2B+D) Integrated Services Digital Network (ISDN) services at 128 kbps. Even at this speed, telecommuters may quickly grow

impatient with slow response times as compared to the throughput of their LANs to which they have grown accustomed. On average, a typical Ethernet user can expect to achieve approximately 1 Mbps on a shared 10Base-T Ethernet LAN and up to 9+Mbps in a full duplex switched

Ethernet environment. In addition, Internet users are also demanding greater access speeds to cope with the various multimedia applications that are continually being developed. Fortunately, the communication industry has recognized the escalating demand.

Cell switching technology, such as Asynchronous Transfer Mode (ATM), was developed in part because of the need to provide a high-speed backbone network for the transport of various types of traffic, including voice, data, image, and video. An ATM network 113 is typically able to provide bandwidths to an ATM user at approximately 1.5 Mbps on a T1 line, 44.7 Mbps on a T3 line, and 155 Mbps over a fiber optic OC-3c line. Consequently, ATM networks are suitable to transport multimedia information.

ATM further provides a mechanism for establishing quality of service (QoS) classes during the virtual channel setup, thereby allotting a predetermined amount of bandwidth to the channel. QoS classes define five broad categories that are outlined, for example, by the ATM Forum's UNI 3.0/3.1 specification. Class 1 specifies performance requirements and indicates that ATM's quality of service should be comparable with the service offered by standard digital connections. Class 2 specifies necessary service levels for packetized video and voice. Class 3 defines requirements for interoperability with other connection-oriented protocols, particularly frame relay. Class 4 specifies interoperability requirements for connectionless protocols, including IP, IPX, and SMDS. Class 5 is effectively a "best effort" attempt at delivery; it is intended for applications that do not require guarantees of service quality.

In conventional data networks, such as the typical Ethernet LAN or X.25 WAN, there are no explicit negotiations between the network and the user specifying the traffic profile and quality of service expected. Rather, the network is expected to provide each user with a "fair share" of the available bandwidth.

However, in an ATM network, fair allocation of bandwidth requires users to adjust their transmission rates according to the feedback from the network. ATM networks carry fixed bandwidth services required for multimedia applications (constant bit rate (CBR) traffic) and guaranteed bandwidth services for high-priority data applications (variable bit rate (VBR) traffic). The remaining bandwidth, not used by guaranteed bandwidth services, must be shared fairly across all users. The ATM Forum refers to services that make use of this otherwise idle bandwidth as available bit rate (ABR) services.

Although these ABR applications must contend for remaining available bandwidth and would not provide specific throughput guarantees, ABR applications still would require fair access to the available bandwidth with a minimum of cell loss. If ABR traffic had no mechanism to

determine if sufficient bandwidth were available to handle the transmission on the network and traffic was simply fed in, network congestion might result in dropped cells, and application traffic might be lost. ABR flow control is an ATM layer service category for which the limiting ATM layer transfer characteristics provided by the network may change after establishing the network connection. A flow control mechanism is specified which supports several types of feedback to control the source rate in response to changing ATM layer transfer characteristics. When the network becomes congested, the end-stations outputting ABR traffic are instructed to reduce their output rate. It is expected that an end-system that adapts its traffic in accordance with the feedback will experience a low cell loss ratio and obtains a fair share of the available bandwidth according to a network-specific allocation policy. Cell delay variation is not controlled in this service, although admitted cells are not delayed unnecessarily.

In this end-to-end rate-based scheme, the source (e.g., a user remote site 103) of a virtual circuit (VC) indicates the desired rate in a resource management cell (RM cell). An RM cell is a standard 53-byte ATM cell used to transmit flow-control information. The RM cell travels on the VC about which it carries information, and is therefore allowed to flow all the way to the destination end-station (e.g., PC 119d). The destination reflects the RM cell, with an indicator to show that the RM cell is now making progress in the reverse direction. The intermediate switches (e.g., switch 109) then identify within the reverse RM cell their respective maximum rates (the explicit rate allocated to the VC). After the source receives the reverse RM cell, the smallest rate identified in the reverse RM cell is then used for subsequent transmissions until a new reverse RM cell is received.

ATM has many recognized advantages and has dominated wide area networks (WANS) as the preferred backbone transport technology. Because of cost and performance factors, ATM faces stiff competition from both switched and shared-media high-speed LAN technologies, including Ethernet, Fast Ethernet, and Gigabit Ethernet. And although ATM typically offers QoS guarantees superior to the prioritization schemes of competing high-speed technologies, many users remain unable to take advantage of these features. If a remote user wishes to obtain the advantages of ATM, one solution would be to acquire an ATM switch on the premises as shown in FIG. 1A. The remote site 103 would need to be equipped with an ATM switch 103a, whereby a PC 103b interfaces the ATM switch 103a via an ATM NIC 103c. In addition, the remote user would have to lease a T1 line or an OC-3c pipe from the Telco. The leased line would terminate in an ATM switch 109 in the CO 105. The CO ATM switch 109 is connected to the ATM network 113. With an ATM connection, the remote user may quickly access multimedia information on the Internet by establishing a virtual channel that would terminate at ATM switch 125 in CO 121. The CO 121 would of course have some means of communication with the ISP 117; typically routers (not shown) are used.

Alternatively, FIG. 1B illustrates an ATM to the desktop solution whereby the xDSL technology is utilized to extend ATM capability remotely. At the customer premises 103, a PC 103b is equipped with an ATM NIC 103c, which is attached to an xDSL modem 103d. In addition, a telephone set 103e is linked to the xDSL modem 103d. The xDSL modem is connected over twisted pair copper wire to the CO 105, terminating at the POTS splitter 117. The POTS splitter 117 separates the data signals originating from the PC 103b

from the voice signals. A xDSL multiplexer (mux) 115 receives the data signals from the POTS splitter and uplinks these signals to the ATM switch 105. Although the solution present above provides a way to deliver ATM capabilities to the desktop, it disadvantageously requires the acquisition of ATM NICs by the remote users, and the xDSL modem has to have a costlier ATM interface.

Despite all the many inherent advantages with ATM, Ethernet-type LANs constitute nearly all of the networking resources of business and residential users. Moreover, these legacy systems are still being enhanced and marketed, e.g., switched Ethernet, switched Fast Ethernet, and switched Gigabit Ethernet are significantly lower cost than their ATM counterparts. ATM technology requires a substantial investment in infrastructure, from cable plant to switches to network interface cards (NICs). This tremendous investment cost can be sustained in the wide area network (WAN) where costs can be spread out. However, in the LAN environment, the investment in infrastructure is typically unsustainable which translates into retention of "legacy" LANs such as Ethernet.

While a number of service providers (e.g., Telcos) employ ATM to establish point-to-point circuits, little has been done to utilize ATM for transporting multimedia information or services to the desktop. This is simply not commercially practical. As previously noted, commercial practicality prohibits such an endeavor. In essence, millions of users would be required to purchase expensive ATM network interface cards, and then possibly add very costly T1, T3, or OC-3c lines. As a result, service providers have not commercially implemented ATM in the delivery of multimedia information to the desktop.

One apparent disadvantage is the inability of current access techniques to adequately provide bandwidth to the end user for access to networking resources.

Another disadvantage is the high cost related to procuring sufficient bandwidth to alleviate long delays caused by low speed connection.

DISCLOSURE OF THE INVENTION

There is a need for an arrangement that enables the high-speed transmission of multimedia information to the desktop.

There is also a need for an arrangement that enables use of an Ethernet-type network interface device to access a high speed cell switching network.

There is also a need for an arrangement that ensures reliable connectivity to networking resources.

These and other needs are attained by the present invention, where a digital subscriber line (DSL) multiplexer employs a proprietary protocol to communicate with an end user DSL modem to monitor the connection.

According to one aspect of the present invention, a multiplexing device for providing high speed network connectivity for the transport of data frames to at least one subscriber, comprises a digital subscriber line (DSL) transceiver interface coupled to plurality of subscriber DSL modems. A packet buffer is coupled to the DSL transceiver interface for storing the data frames. The packet buffer transfers the data frames over a packet buffer bus. An Ethernet-type port is coupled to the packet buffer for uplinking to an external switch. A central processing unit (CPU) is coupled to the packet buffer for executing code to send and receive proprietary protocol packets for communicating management information associated with the plurality of

subscriber DSL modems. The management information indicates presence of an Ethernet-type port, current line speed, and media access control (MAC) address.

Additional advantages and novel features of the invention will be set forth in the description which follows, and in part will become apparent to those skilled in the art upon examination of the following or may be learned by practice of the invention. The advantages of the invention may be realized and attained by means of the instrumentalities and combinations particularly pointed out in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

Reference is made to the attached drawings, wherein elements having the same reference numeral designations represent like elements throughout and wherein:

FIGS. 1A and 1B are graphic representations of a prior art networks and the access methods;

FIGS. 2A and 2B are block diagrams depicting detailed aspects of systems incorporating a multiplexer configured in accordance with the present invention;

FIG. 3 is a block diagram of the hardware architecture of a DSL multiplexer in accord with one embodiment of the present invention;

FIG. 4 is a block diagram illustrating a boot code architecture;

FIG. 5 is a block diagram illustrating a runtime code architecture;

FIG. 6 is a block diagram of a standard Ethernet type 2 packet format;

FIG. 7 is a block diagram of a data field of a proprietary protocol PDU; and

FIG. 8 is a block diagram of a generic PDU format for exchanging information in accordance with the present invention.

BEST MODE FOR CARRYING OUT THE INVENTION

The present invention provides a high speed connection to the end user through a digital subscriber line multiplexer. This retains the traditional low cost and low complexity associated with Ethernet technology while achieving QoS. NIC but achieves ATM capability over Ethernet through use of an Ethernet edge switch which employs a multi-processor architecture to interface an Ethernet environment with an ATM infrastructure.

FIG. 2A illustrates an embodiment of the present invention taking advantage of the existing network media created by the telephone industry which implemented a vast network of copper twisted pair wiring to interconnect homes and businesses domestically and abroad. In FIG. 2A, a customer premises 200 is shown as comprising an end-station 210, such as a desktop computer residing in a home or business. Typically, such end-stations are either stand-alone desktop stations, or are already connected to a collocated local area network (LAN).

A variety of LAN technologies exist, but the large majority of LANs conform to the IEEE standard 802.3, which defines Ethernet standards. Various types of Ethernet systems exist such as switched Ethernet, Fast Ethernet, and Gigabit Ethernet. The end-station 210 is equipped with an Ethernet NIC 212 residing in a personal computer with a host processor 214. Ethernet NIC 214 is connected to a high-speed digital subscriber line (DSL) modem 220, which

interfaces with a telephone line 222 via a CP (customer premise) POTS (plain old telephone service) splitter 221. The telephone line is a twisted pair copper wire, which the conventional customer premises telephone 224 uses to connect with a telephone communication facility 240. A telephone communication facility 240 or end office is shown in FIG. 2A as a communications facility 240; however, any communication facility can be used (e.g., a wire closet in a separate building).

High-speed communication to remote users depends largely on the method of access to the networking infrastructure. Most users cannot bear the cost of leasing expensive outside lines that are needed to provide high speed communication to the Internet or to their offices. The disclosed embodiment overcomes this dilemma by employing a high-speed, low cost subscriber interface that takes advantage of the legacy outside cable plant, such as standard twisted copper pair wiring and coaxial cables.

One embodiment, shown in FIG. 2A, utilizes digital subscriber line (DSL) technology to deliver the high bandwidth that the remote users demand. Because traditional copper cabling is used, the remote users do not have to upgrade their current physical connection—their POTS line is sufficient. Because the outside plant need not be revamped, telephone companies (Telcos) can readily implement DSL services. The DSL modem 220 acts as the network access device to the communication facility 240. A DSL multiplexer 252 provides termination of the DSL modem connection within communications facility 240. DSL technology is categorized by the downstream and upstream bandwidths. The present invention could be applied to any of the various forms of DSL technology. One variety, commonly employed, Rate Adaptive DSL or RADSL, involves a rate negotiation between the customer premise DSL modem 220 and the Telco CO modem located within DSL MUX 252 which takes into account distance and line quality issues yielding the maximum available rate for the line conditions encountered. RADSL supports both Asymmetric DSL or ADSL, with a maximum downstream rate of 7.62 Mbps and a maximum upstream rate of 1.1 Mbps, which is ideal for very high speed Internet access and video-on-demand applications. ADSL services can be delivered up to 18,000 feet from the communication facility 240 over a single copper twisted pair. RADSL also supports Symmetric DSL or SDSL, with a maximum bi-directional rate of about 1.1 Mbps, which is ideal for very high quality video-conferencing and remote LAN access. Another type of DSL technology is known as high-bit-rate digital subscriber line (HDSL), which provides a symmetric channel, delivering T1 rates (1.544 Mbps) in both directions. HDSL has a distance limitation of about 12,000 feet without repeaters. Telcos have traditionally used HDSL to provide local access to T1 services. HDSL is already widely deployed within the Telco market as a low cost T-1 replacement. VDSL or Very high bit-rate DSL requires a fiber-to-the-curb local loop infrastructure, with asymmetric speeds up to 52 Mbps. Other flavors of DSL (i.e., sometimes generically denoted xDSL) are characterized by whether the service is asymmetric or symmetric and the bandwidth allocations for the upstream and downstream transmissions.

The communication facility 240 comprises a plain old telephone service (POTS) splitter 242 which receives the information transmitted across the twisted pair line 222 and "splits" the low frequencies, which carry voice signals, from the high frequencies, which carry data signals. Essentially, the POTS splitter is a passband filter, whereby the low frequency information is carried by a voice line 224 to a

voice switch 246 and ultimately to a public switched telephone network (PSTN) 248. The voice line 224, voice switch 246 and PSTN 248 are each conventional, and are therefore not explained further so as not to detract from the focus of the disclosure of the present invention.

The data information, which is modulated using high frequency signals, is transmitted over a twisted pair cable 250 to a POTS splitter 242. The POTS splitter 242 then passes the high frequency signals to a DSL multiplexer (DSL MUX) 252. The DSL MUX serves as the DSL modem termination point for numerous end users with DSL modems. The DSL MUX 252 aggregates all the DSL traffic and passes the multimedia information to the multimedia switch 260. The traffic can be of any data type including multimedia graphics, video, image, audio, and text. Various embodiments of the DSL MUX 252 can be employed, ranging from 24 line stackable modules through the traditional high density chassis based approach. Various line codes can be supported within the DSL modems, including Carrierless Amplitude Phase (CAP) modulation, Discrete Multi-Tone (DMT) modulation, Quadrature Amplitude Modulation (QAM), as well as others. Multimedia switch 260 is primarily an edge device that is connected to an ATM network 270 on which a conventional multimedia server (not shown) may be linked. The ATM network 270 thus represents a fast and efficient delivery system for multimedia applications to which the end user desires access. The multimedia switch 260 communicates with the CO DSL MUX 252 relative to traffic information, in order to minimize congestion. Traditionally, end user access to an ATM network has been through a router. Since the end-station 210 houses an Ethernet NIC 214, connection to ATM network 270 proves difficult without the system of the present invention, which allows information residing on an ATM network to be transferred to an Ethernet end-station while still retaining all the multimedia benefits of ATM, including QOS and ABR/ER flow control. An advantage associated with a DSL implementation is that the personal computer is constantly connected, much like a typical Ethernet LAN connection. That is, communication sessions are not initiated through a dial-up procedure.

FIG. 2B illustrates another embodiment of the present invention where DSL Mux 252 is able to connect to a plurality of DSL modems 220. This configuration enables many customer premise modems 220, for example twenty four, to be pooled together. Further, the multimedia switch 260 can accommodate multiple DSL multiplexers 252, thereby achieving a higher density of subscribers. The uplink from the DSL multiplexer 252 to the multimedia switch 260 is typically a fast Ethernet connection.

Turning now to FIG. 3, DSL multiplexer 252 comprises a number of an ADSL modem devices based on CAP (Carrier-less Amplitude/Phase) modulation technology, and is located within communication facility 240. The device will be used for Internet access and corporate private network use. DSL multiplexer 252 has station 6modems 300, which can be connected to 24 subscriber lines (SL) (e.g., 222) having customer premise modems, for example, DSL modem 220. DSL multiplexer 252 also has two 100Base-T ports 302 to communicate with multimedia switch 260. In such a configuration, multimedia switch 260 will have up to 24 end stations, such as end station 210, connected to one 100Base-T port. One or more of these end-stations can be CIF end stations. In order to deliver QoS to these end stations, some information has to be passed between DSL multiplexer 252 and multimedia switch 260. A specific protocol is defined for this purpose. DSL multiplexer 252

can implement, for example, the Intel 80960HD (i960HD) 32 bit RISC processor running at 33 MHz.

Two 100BaseT ports 302 operate as one primary port and one standby port. The standby port takes over the functions of the primary port when the primary port fails and vice versa. At any time, only one port will be active. 100BaseT port 302 is implemented in one embodiment using Intel's 82557 MAC and external transceivers. 100baseT port 302 supports full duplex operation. Logic is provided to drive the "link/activity" and "port disabled" LEDs. The physical connector is a standard RJ45 jack.

Two DRAM based memory banks, i.e. a packet buffer 304 and a local buffer 306, are implemented in one embodiment. Packet buffer 304 is used for the DSL data and local buffer 306 is used as code memory and workspace. Fast page EDO (Extended Data Out) DRAMs are used in an embodiment for both memory banks, since they can work with zero wait states in a 33 MHz system without having to introduce memory interleaving.

Remote access controller(RAC) 308 provides the MAC functions for the 10BaseT port used for SNMP traffic. This port can also be used as a high-speed debug port during software testing. External 10Base-T transceiver is provided. Logic is provided to drive the "link/activity" and "port disabled" LEDs. The physical connector is a standard 8 pin RJ45 jack. The RAC chip 308, for example the Galileo Technology (GT96010), is used to interface the ADSL serial data (to/from the GTI modem devices) to packet buffer 304. This chip has six multi-protocol serial channels and an Ethernet port. It has a DMA engine which is compatible with, for example, i960JX processor. Some glue logic is required to interface this chip to i960HD. Each GT96010 can currently support three ADSL channels at maximum downstream data rate of 8 Mbps and an upstream rate of 1 Mbps. The two 100baseT uplink ports 302 can be implemented using the Intel 82557 MAC chip. The 100baseT ports 302 support full duplex operation.

A local bus 310 will support a CPU 312, local DRAM 306, a flash memory 314, an EPROM 316, a NVSRAM 40 (non-volatile static random access memory) 318, control ports and status ports 320, a DUART 322 and a set of buffers 324 which connects to the other buses. Flash memory 314 (512Kx8) is used for the storage of firmware. On power-up, the firmware will be downloaded from Flash 314 into local DRAM 306 (4MB) for execution. The boot code resides in EPROM 316. DUART 322 provides (a) the PPP SNMP link (also used by a local diagnostic and configuration utility) and (b) a debug port which is used during the development cycle for testing and debugging the board. All modem transceivers 300 sit on an 8 bit wide modem bus 325. Crosspoint switches 326 sit on an 8 bit wide switch bus 327.

Packet buffer 304 (4 MB) is used to store the frames from the DSL multiplexer 252 interfaces (100BaseT and 10BaseT). The 100BaseT MAC chips can be a PCI device 55 and can connect to packet buffer 304 through a PCI bridge 330. CPU 312 accesses the packet memory through buffer 304. Remote access controller (RAC) chip 308 is used to provide the 10BaseT interface for SNMP traffic. The six modem transceivers 300 are interfaced to the packet buffer 304 using two remote access controllers 309. All the above masters on a packet bus 332 operate with multiplexed address and data. A common address latch 334 is used to route the addresses to packet buffer 304. Arbitration among the contenders for packet buffer 304 is done in round robin fashion.

Two serial ports 323, operating at a maximum speed of 38.4 Kbps will be provided by using DUART chip 322. One

serial port will be used for SNMP traffic over a PPP (point-to-point protocol) link. The same serial port can also be used for running the local DSL multiplexer 252 diagnostics utility. This port will also support a null modem connection (auto-answer only). The second serial port will be used only during development for debugging the card. The physical connector for each port is a 9 pin D-Sub receptacle.

The transceiver interface will be based on a RADSL chipset implemented as DSL modem 220. Six modem chipsets will be used to connect to 24 subscriber lines through a cross point switch. The GT96010 multi-channel serial controllers will be used to interface the transceivers to packet buffer 304. Each controller provides three ADSL interfaces and a DMA engine supporting data bursts to the memory. DSL multiplexer 252 will use an external POTS splitter 242. One 50 pin champ connector will be used for the incoming DSL lines from POTS splitter 242.

The two control ports 320 use 32 bits, of which some bits are used as individual resets to the various controllers. 48 bits drive DSL status LEDs, 2 bits drive health/diagnostics LEDs, and 2 bits drive power-supply status LEDs. Three 8-bit ports control the power to line-drivers (SLIDES) 336. A fourth 8-bit port provides individual resets to modem transceivers 300. One 32-bit status port is provided to read in the status of the interrupt lines and of the power supply outputs. Another 8-bit port is used to read the interrupt status of modem transceivers 300.

DSL multiplexer 252 supports either (a) DC inputs (with internal DC-DC converters), for the Telco environment or (b) AC inputs (with internal SMPS) for corporate environments. In the Telco case, dual redundant DC-DC converters and in the corporate case, dual redundant SMPSs will be mounted inside the unit. Each of these will provide 5V, 3.3V and +/-15V with sufficient current to drive the electronics. The power supplies used will be having the protection diodes and the corresponding outputs will be shorted together. A faulty power supply can be without shutting down DSL multiplexer 252.

Since DSL multiplexer 252 is a store-and-forward device, data forwarding module 510 receives data on the ADSL modem lines which is multiplexed and sent over the 100 Mbps Ethernet port 302 (FIG. 3). The source MAC address of these packets are stored in an address table that pairs subscriber line number with MAC addresses. This address table will be needed when a frame is received over 100 Mbps Ethernet port 302 in order to perform data demultiplexing. The destination MAC address in the frame is used to index the table and identify the subscriber line to which the data has to be sent. The software in DSL multiplexer 252 comprises two components: Boot code and Runtime code. The Boot code resides in EPROM 316 and is invoked at power-on. It initializes the processor and executes the power-on diagnostics code. If fatal errors are seen, the boot code will blink the health LED and display the error code on the DSL status LEDs, and do nothing further. If the tests pass, and if Runtime code is present in the flash EPROM 314, the boot code will then decompress and move the runtime code from flash memory 314 to local memory 306 and transfer control to the runtime code. If the runtime code is not present in flash memory 314, the code is downloaded from a TFTP server either through the serial link, the 10Base-T link or through the 100Base-T link. The boot code also contains a debugger, configuration utility and diagnostic utility over the RS232 port.

The architecture of the runtime code is shown in FIG. 5. A supervisory module 500 has an initialization module and

a scheduler. The initialization module does all the startup initialization. The scheduler executes the different tasks defined in the software in a round-robin fashion. The main functions of the runtime code include: forwarding the data from the ADSL lines to 100Base-T port 302, demultiplexing and forwarding data from 100Base-T port 302 to the ADSL lines, controlling the operation of the modem transceivers 300, and sharing the modem transceivers 300 among the 24 subscriber lines. The Runtime code also includes a SNMPv2 based agent. The Boot code architecture is shown in FIG. 4. A processor initialization module 400 initializes the serial port, and the processor control structures on power-up/reset. Then, the power-on diagnostics are run.

A power-on diagnostics module 402 performs tests on the various hardware blocks. The result of the diagnostics is stored in NVSRAM 318. If the tests are passed, the green "Health" LED will blink at a slow rate of 1 second on and 1 second off. A supervisory module 404 downloads the runtime code from flash memory 314 to local memory 306 and executes it. If the runtime code is not present in flash memory 314, the code is downloaded from the TFTP server. supervisory module 404 is a simple control loop with tasks executed in a round-robin fashion.

A TFTP module 406 is used for downloading the runtime code from a TFTP server and uploading the code to a server if needed. A UDP/IP stack 408 the protocol modules IP and UDP. A PPP module 410 implements the PPP protocol. A flash utilities 412 stores runtime code in flash memory 314 in compressed form, and contains all the functions related to flash memory 314 read/write routines and compression-decompression algorithms. A 100Base-T driver 414 and a 10Base-T driver 416 are each used during execution of Boot code for downloading the runtime file from the TFTP server, if required. Driver 414 controls the operation of sending Ethernet packets, receiving Ethernet packets and allocation and management of receive and transmit queues. A serial driver 416 is used during execution of Boot code for printing debug messages and also for downloading the runtime code during testing. The serial driver controls the operation of the serial interface chip. A debugger 418 comprises typical debugger functionality, and also includes a menu driven interface for local diagnostics. A configuration utility 420 allows the user to configure the DSL multiplexer 252. The user can access this utility through the serial port. System utilities 422 include a timer and memory buffer functions.

Initially, when the runtime code is launched, a modem pool module 502 connects a CO modem to each of the subscriber lines in turn. For each connection, the line parameters such as maximum upstream and downstream speeds, receiver equalizer coefficients, etc are retrieved from the modem and stored in a database in memory. This database will be used later to "reestablish" a connection. Modem-pool module 502 "scans" all unconnected lines, in rotation, using free modems (if no free modem exists, it will not scan). If more than one free modem are available, the scanning will be done on many lines in parallel. After connecting the CO modem to a particular line, the line parameters of that line will be stored to the CO modem and a local warm startup command (ATT_LCL_WARM_REQ) is given with a max time-out of 3 second. The CP modem 220 will detect this command and will respond to the warm startup and both CO modem and CP modem will come to the data mode. If the CP modem 220 has no data to transfer it will issue a local standby command (ATT_LCL_STBY_REQ) immediately. When the inactivity is detected by CP modem 220 at the Ethernet port (no data for more than TBD minutes, may be one minute) the CP modem 220 issues

a local standby command. The CO modem in the DSL multiplexer 252 will detect this condition and will go to the standby mode and release the modem to the free pool.

A RADSL transceiver driver 504 controls the operation of the ADSL transceivers present on the card. RADSL transceiver driver 504 handles start-up of the CO modems, and setting up the parameters and reading status of the CO modems. An ADSL driver 506 controls the operation of RACs 308, 309 (GT-96010's) interfaced with transceivers 300. ADSL driver 506 sends and receives packets over the ADSL line using the HDLC-like framing mode. The parameters downstream baud rate, constellation, noise margin, transmit power level etc. are stored in the NVRAM 318.

100 Mbps Ethernet port driver 508 provides the up-link to the data network. 100 Mbps Ethernet port driver 508 controls the operation of sending Ethernet packets, receiving Ethernet packets and allocation of receive and transmit queues. This also handles the switching to the secondary link when the primary link fails.

ASNMP and MIBs module 512 implements the SNMPv2 protocol and thus provides management and statistics collection through both in-band and out-of-band SNMP links.

A UDP/IP stack are implemented in a UDP/IP module 516. Minimal routing functions are implemented in the IP layer, to allow various network management configurations. DSL multiplexer 252 will route SNMP packets destined for other DSL multiplexers 252 as follows: (a) packets received on the serial port will be sent either to the 10Base-T port or to the 100Base-T port depending on the routing table, (b) packets received on the 10Base-T port will be sent on the 100Base-T port. (response packets follow the reverse path.) Routing tables will be manually entered by the operator via the configuration utility; they will not be dynamically updated by the IP module.

10Base-T driver 518 controls the operation of the 10Base-T Ethernet interface. 10Base-T driver 518 performs the functions of: sending Ethernet packets, receiving Ethernet packets, and allocation and management of receive and transmit queues.

PPP module 520 implements the point to point protocol (PPP) and interacts with a serial driver 522 which controls the operation of the serial interface chip

A proprietary protocol handler 524 is used for sending/receiving proprietary protocol packets on the ADSL interface for managing CP modem 220.

The following parameters have to be passed from DSL multiplexer 252 to multimedia switch 260: presence of DSL multiplexer 252 on a Fast Ethernet port, current ADSL link speed on each one of the CP modems 220 connected to DSL multiplexer 252, MAC address of end station 210 connected to each CP modems 220.

DSL multiplexer 252 sends unsolicited Ethernet packets to multimedia switch 260 at regular intervals. At present the time interval between two successive packets is fixed, for example at 300 seconds. A MIB variable will be provided to configure this parameter. The interval is kept large enough so as not to consume any significant bandwidth on the 100Base-T link. Multimedia switch 260 can request DSL multiplexer 252 for any of the above parameters at any time. On receiving such a request, DSL multiplexer 252 responds with a packet containing the value of the requested parameter.

FIG. 6 is a graphical representation of a standard Ethernet type II (DIX) packet format 600. A type field 602 in the Ethernet packet is a value greater than or equal to 0800h.

This gives the SAP number of the protocol destined to receive this packet. Multimedia switch 260 and DSL multiplexer 252 use a proprietary type value of FFF0h for sending their protocol packets.

5 A data field 700 of the protocol PDU is graphically represented in FIG. 7. A command field (CM) 702 comprises one byte, a total length (TL) field 704 comprise two bytes and includes the command field 702. A sequence number (SN) field 706 comprises one byte. A type (T) field 708, a length (L) field 710 and a value (V) field 712 are also provided.

10 The command field 702 can have the following values:

15	01h	Request parameter (REQ)
	02h	Report Parameter (REP)
	03h	Unsolicited parameter update (UPU)
	04h	Keep-alive (KA)
	0ffh	Acknowledgement (ACK)

20 More enumerations will be added as and when necessary.

The values of command field 702 are used in a protocol administered between DSL multiplexer 252 and multimedia switch 260. After start up and initialization, DSL multiplexer

25 252 starts sending keep-alive (KA) packets on its 100Base-T port 302 once in 5 minutes. This packet will have CM=4, LN=4 and no TLV fields 708, 710, 712. SN field 706 will start from 0 and gets incremented for every subsequent packet. This field wraps around to 0 after ff.

30 If multimedia switch 260 receives this packet, it sends an acknowledgement (ACK) packet to DSL multiplexer 252. This packet will have CM=ff, LN=5, SN same as the value in the received KA packet and no TLV fields. One byte, immediately after the SN field 706 contains the value of CM field 702 that was in the received KA packet.

35 To start with, DSL multiplexer 252 initializes its 5800_discovered flag to FALSE. When DSL multiplexer 252 gets an ACK for the keep-alive packet, it sets 5800-discovered flag to TRUE. Further on, if it does not get an ACK packet

40 for 2 successive KA packets, it resets 5800-discovered flag to FALSE. However, DSL multiplexer 252 will continue to send KA packets.

Multimedia switch 260 initializes its per-port 5300_discovered flag to FALSE. If multimedia switch 260 receives a KA packet on a port, the 5300_discovered flag for that port is set to TRUE. Further on, if multimedia switch 260 does not receive a KA packet on a port for 2 minutes, the 5300_discovered flag for that port is set to FALSE.

If the 5800_discovered flag is TRUE, DSL multiplexer 252 sends an unsolicited parameter update (UPU) packet when it detects any change in the exchanged configuration parameters. The UPU packet will also be sent with all the parameters when the 5800_discovered flag changes from FALSE to TRUE.

55 The UPU packet contains variable number of the following TLV fields: (a) a MAC address having a type field of one byte, a length field of eight bytes, a value field of six bytes of DSL multiplexer MAC address; (b) link speed having a type field of two bytes, a length field of five bytes, and a

60 values field of 3 bytes—1 SL number, 2 link speed upstream, 3 link speed downstream; and (c) a port MAC address having a type field of three bytes, a length field of nine bytes, and a value field of seven bytes. The seven bytes of the value field correspond to SL number (1) and MAC address (2-7).

65 When multimedia switch 260 receives a UPU packet, it updates its database with the information contained in the packet. Then it sends an ACK packet for the UPU packet.

Multimedia switch 260 can any time send a Request Parameter (REQ) packet to DSL multiplexer 252 on one of its ports. This packet will have the following format:

CM	01h
T	Type of the requested parameter
L	03h
V	SL number

DSL multiplexer 252 can request multiple parameters in a single PDU.

When DSL multiplexer 252 gets a REQ packet from multimedia switch 260, it fills the values of all requested parameters in the REQ PDU and sends a Report Parameter (REP) PDU to multimedia switch 260. The REP PDU fields are given below:

CM	02h
SN	Same as in the REQ PDU
T	Type of the parameter
L	Length of the TLV field
V	SL number and value.

FIG. 8 illustrates a generic PDU format 800 for exchanging proprietary protocol between modem 220 and DSL multiplexer 252. Generic PDU format 800 includes an Ethernet header 802, a DSL multiplexer header 804, data field 806, and two bytes of check sum 808. An Ethernet header 802 comprises a source MAC which contains a 10BaseT MAC address of CP Modem 220 when a PDU is originating from the modem 220. When the PDU originates from the CO DSL multiplexer 252, this field contains its 10BaseT/100BaseT Ethernet Port's MAC address.

Ethernet header 802 also includes a destination MAC which contains a broadcast MAC address irrespective of whether the packet is originating from modem 220 or CO DSL multiplexer 252. Ethernet header 802 has a type/length which should contain 0 to uniquely differentiate between a Standard Ethernet packet and a DSL multiplexer 252 packet on the ADSL link. Standard Ethernet Packets will have a non-zero value for this field which signifies that the packet format is either DIX or IEEE.

The DSL multiplexer 252 can be managed through the console port, 10Base-T port or 100Base-T port. The DSL multiplexer 252 will provide SNMP proxy functionality for allowing management of the CP modems connected to it. A proprietary protocol will be used between the CP & DSL multiplexer 252 to transact management information between them. This protocol will be transparent to the NMS. The DSL multiplexer 252 maintains the latest management information from all the CPs connected to it. This information will be obtained by the DSL multiplexer 252 through periodic polling of the CP modems 220 which are currently on-line. The DSL multiplexer 252 routes the management traffic among the following 3 ports: 10Base-T, 100Base-T and Console Port. Hence, each of these ports has to be allocated an IP address (belonging to different subnets). This will enable an NMS present on any of these 3 subnets to manage other DSL multiplexer 252s within the Telco. The DSL multiplexer 252 will use statically configured routing table entries for routing.

The Ethernet switch 260 as detailed above works in conjunction with the end-user-workstation's software to quickly deliver multimedia information while ensuring an end-to-end negotiated quality of service that is free from

delay inducing congestion. The end-station executes a shim software. The shim comprises a protocol combination, or other suitable combination of protocols, to allow the implementation of CIF technology to bring native ATM services to desktops that are equipped with legacy Ethernet or Token Ring NICs by encapsulating cells into frames. CIF can also be viewed as the inverse of ATM LAN Emulation (LANE). LANE provides a way for legacy LAN media access controller-layer protocols like Ethernet and Token Ring, and all higher-layer protocols and applications, to access work transparently across an ATM network. LANE retains all Ethernet and Token Ring drivers and adapters; no modifications need to be made to Ethernet or Token Ring end stations. In other words, CIF emulates ATM services over frame-based LANs. CIF uses software at the workstation without requiring the procurement of a new NIC to support quality of service scheduling and ABRIER flow control.

To achieve end-to-end quality of service, the shim resides as a layer in end station to provide encapsulation of cells within Ethernet frames in the desktop for transport to the data network. Shim supports multiple queues, a scheduler (not shown), the ER flow control, and header adjustment. Shim comprises an ATM Adaptation Layer (AAL) which is the standards layer that allows multiple applications to have data converted to and from the ATM cell. AAL is protocol used that translates higher layer services into the size and format of an ATM cell. The CIF shim layer also includes a traffic management (TM) component that sets forth the congestion control requirements. The TM component (not shown) can be implemented as TM 4.0. The ATM Forum has developed a complete 4.0 protocol suite that includes UNI signaling 4.0 which allows signaling of bandwidth and delay requirements for QoS; whereby, TM 4.0 which specifies explicit rate flow control and QoS functions.

CIF shim layer also includes a frame segmentation and reassembly (SAR) sublayer (not shown) which converts protocol data units (PDUs) into appropriate lengths and formats them to fit the payload of an ATM cell. At the destination end station, SAR extracts the payloads for the cells and converts them back into PDUs which can be used by applications higher up the protocol stack. The shim adds the CIF header to packets before they are transmitted, and removes the header when they are received. The shim manages the message queues by queuing outgoing data into multiple queues for QoS management. Shim also processes the RM cells for explicit rate flow control using the ABR flow control and allows ATM signaling software to run both native ATM application as well as standard IP applications.

End station further comprises a device driver and a Network Device Interface Specification (NDIS) layer 609 located above the CIF shim layer 611. The end station 714 includes Internet Protocol (IP) layer 607b which supports classical IP, LANE and MPOA for the interworking of dissimilar computers across a network. IP layer 607b is a connectionless protocol that operates at the network layer (layer 3) of the OSI model. Winsock 2.0 603 is the application program interface (API) layer, which enables developers to take advantage of ATM's QoS and traffic management features. Application layer 601 can accommodate traditional as well as native ATM applications. Native ATM applications can be readily created with Winsock 2.0 API 603.

The shim arrangement guarantees that the services negotiated by the native ATM applications for the VCs are not arbitrarily disrupted by the traffic generated by the legacy applications. Forcing both the ATM and the legacy protocol traffic to go through CIF shim allows CIF shim to manage

the transmission of all traffic according to the QoS specified for each traffic stream. To support the migration of legacy applications, the CIF AD forwards CIF traffic from the conventional LAN onto the ATM infrastructure for delivery to an ATM attached end station or to another CIF AD. The CIF ES is also required to run LANE, MPOA (Multiprotocol Over ATM), or Classical IP protocols. Network data from a legacy application is first handled by the legacy protocols (e.g., TCP/IP), and then turned into ATM traffic by LANE, MPOA, or Classical IP. The CIF ES function encapsulates the individual cells into CIF frames before data is finally transmitted on the wire to the CIF AD.

The enhancements in the network as discussed above can be implemented if the end-user is shielded from any bottlenecks that will negate such enhancements; the bottleneck typically exists at the access link. Thus, the method of access needs to be made fast and reliable.

While this invention has been described in connection with what is presently considered to be the most practical and preferred embodiments, it is to be understood that the invention is not limited to the disclosed embodiments, but, on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims.

What is claimed is:

1. A multiplexing device for providing high speed network connectivity for the transport of data frames to at least one subscriber, comprising:

a digital subscriber line (DSL) transceiver interface coupled to a plurality of subscriber DSL modems;

a packet buffer coupled to the DSL transceiver interface for storing the data frames, the packet buffer transferring the data frames over a packet buffer bus;

an Ethernet-type port coupled to the packet buffer for uplinking to an external switch; and

a central processing unit (CPU) coupled to the packet buffer for executing code to send and receive proprietary protocol packets for communicating management information associated with the plurality of subscriber DSL modems, the management information indicating presence of an Ethernet-type port, current line speed, and media access control (MAC) address;

wherein the proprietary protocol packets comprise:

a line identification field that identifies a subscriber line associated with one of the subscriber DSL modems;

a sequence identification field that is used for synchronization between the data frames;

a data length field that specifies the length of a data payload within the data frames; and

a checksum field for error detection.

2. The multiplexing device as in claim 1, further wherein an external POTS splitter supplies DSL signals to the DSL transceiver interface.

3. The multiplexing device as in claim 1, further wherein the subscriber DSL modem and the DSL transceiver interface utilize a rate adaptive DSL scheme.

4. The multiplexing device as in claim 1, wherein the Ethernet-type port is a 100Base-T port, wherein the Ethernet-type port supports full duplex operation.

5. The multiplexing device as in claim 1, further comprising a cross-point switch coupled to the DSL transceiver interface for modem pooling, wherein the DSL transceiver interface comprises six DSL modems to support at least twenty-four subscribers.

6. The multiplexing device as in claim 1, wherein the management information comprises lower and upper Service Access Point (SAP) information.

7. The multiplexing device as in claim 1, wherein the management information comprises a file download module for retrieving a runtime file from the multiplexer, wherein the runtime file, when executing, is configured to forward data from the DSL transceiver interface to the Ethernet-type port, forward data from the Ethernet-type port to the DSL transceiver interface, control operation of the DSL transceiver interface, and share the DSL transceiver interface among the plurality of subscriber DSL modems.

8. The multiplexing device as in claim 1, wherein the DSL transceiver, the packet buffer, the Ethernet-type port, and the CPU are housed in a stackable chassis.

9. A multiplexing device for providing high speed network connectivity for the transport of data frames to at least one subscriber, comprising:

a digital subscriber line (DSL) transceiver interface coupled to a plurality of subscriber DSL modems;

a packet buffer coupled to the DSL transceiver interface for storing the data frames, the packet buffer transferring the data frames over a packet buffer bus;

an Ethernet-type port coupled to the packet buffer for uplinking to an external switch; and

a central processing unit (CPU) coupled to the packet buffer for executing code to send and receive proprietary protocol packets for communicating management information associated with the plurality of subscriber DSL modems, the management information indicating presence of an Ethernet-type port, current line speed, and media access control (MAC) address;

wherein the management information comprises a file download module for retrieving a runtime file from the multiplexer, wherein the runtime file, when executing, is configured to forward data from the DSL transceiver interface to the Ethernet-type port, forward data from the Ethernet-type port to the DSL transceiver interface, control operation of the DSL transceiver interface, and share the DSL transceiver interface among the plurality of subscriber DSL modems.

* * * * *